IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION

DAEDALUS BLUE LLC, Plaintiff, v. MICROSOFT CORPORATION, Defendant.	Case No. 6:20-cv-01152-ADA

DECLARATION OF MARKUS JAKOBSSON, PH.D., IN SUPPORT OF DEFENDANTS' CLAIM CONSTRUCTION BRIEF

TABLE OF CONTENTS

1. I, Markus Jakobsson, Ph.D., hereby declare as follows:

I. INTRODUCTION

2. My name is Markus Jakobsson. My residence and place of business is at 118 Ramona Rd, Portola Valley, CA. I am over the age of eighteen and a citizen of the United States. I have been retained by Defendant Microsoft Corporation ("Microsoft") to provide my opinions regarding U.S. Patent Nos. 7,437,730 (Def's Br. Ex. 3, the "'730 Patent") and 8,572,612 (Def's Br. Ex. 6, the "'612 Patent") (collectively, the "Asserted Patents"). Specifically, I have been asked to consider claim construction issues pertaining to the Asserted Patents in view of the understanding of a person of ordinary skill in the art ("POSITA"). If called upon to do so, I would testify competently to the facts and opinions herein.

3. I am being compensated for my time at my standard consulting rate. I am also being reimbursed for expenses that I incur during the course of this work. My compensation is not contingent upon the results of my study, the substance of my opinions, or the outcome of any proceeding involving the Asserted Patents. I have no financial interest in the outcome of this matter.

4. My opinions are based on my years of education, research, and experience, as well as my investigation and study of relevant materials, including those cited herein.

5. I may rely upon these materials, my knowledge and experience, and/or additional materials to rebut arguments raised by Plaintiff. Further, I may also consider additional documents and information in forming any necessary opinions, including documents that may not yet have been provided to me.

6. My analysis of the materials produced in this proceeding is ongoing and I will continue to review any new material as it is provided. This declaration represents only those

opinions I have formed to date. I reserve the right to revise, supplement, and/or amend my opinions stated herein based on new information and on my continuing analysis of the materials already provided.

II. QUALIFICATIONS

7. Exhibit 1 is a copy of my current curriculum vitae ("CV").

8. I received a Master of Science degree in Computer Engineering from the Lund Institute of Technology in Sweden in 1993, a Master of Science degree in Computer Science from the University of California at San Diego in 1994, and a Ph.D. in Computer Science from the University of California at San Diego in 1997, specializing in Cryptography. During and after my Ph.D. studies, I was also a researcher at the San Diego Supercomputer Center and General Atomics, where I did research on authentication, privacy, and security. Much of my work was on distributed protocols to enable security and privacy, and my Ph.D. thesis described a distributed electronic payment system. Whereas the term "cloud computing" had not been popularized at this point, many of the principles I studied later became critical aspects of cloud computing and virtualization.

9. Since earning my doctorate degree over twenty-four years ago, I have been an employee, entrepreneur, and consultant in the computer systems, security, and cloud computing industry. I have worked as an engineer at a number of leading companies in the computer industry, including PayPal, Qualcomm, Bell Labs, and LifeLock. I have also taught and performed research at several prestigious universities, including New York University (NYU) and Indiana University (IU).

10. From 1997 to 2001, I was a Member of the Technical Staff at Bell Labs, where I did research on authentication, privacy, multi-party computation, contract exchange, digital commerce including crypto payments, and fraud detection and prevention. Example publications

of mine from this time period include "Secure distributed computation in cryptographic applications" (U.S. Patent No. 6,950,937, 2001 priority date) and "Secure server-aided signature generation" (International Workshop on Public Key Cryptography PKC 2001: Public Key Cryptography 383-401), both of which relate to virtualization.

11. From 2001 to 2004, I was a Principal Research Scientist at RSA Labs, where I worked on predicting future fraud scenarios in commerce and authentication and developed solutions to those problems. Much of my work related to the development of distributed protocols. During that time, I predicted the rise of what later became known as phishing. I also laid the groundwork for what I termed "proof of work," and described how this could be applied in the context of distributed computations such as mining of crypto payments; this work later came to influence the development of BitCoin. I was also an Adjunct Associate Professor in the Computer Science department at New York University from 2002 to 2004, where I taught cryptographic protocols, with an emphasis on distributed computation.

12. From 2004 to 2016, I held a faculty position at Indiana University at Bloomington, first as an Associate Professor of Computer Science, Associate Professor of Informatics, Associate Professor of Cognitive Science, and Associate Director of the Center for Applied Cybersecurity Research (CACR) from 2004 to 2008; and then as an Adjunct Associate Professor from 2008 to 2016. I was the most senior security researcher at Indiana University, where I built a research group focused on online fraud and countermeasures, resulting in over 50 publications and two books. One of these books, "Crimeware: Understanding New Attacks and Defenses" (Wiley, 2008), described the benefits of virtualization in the context of Internet security.

13. While a professor at Indiana University, I was also employed by Xerox PARC,

PayPal, and Qualcomm to provide thought leadership to their security groups. I was a Principal Scientist at Xerox PARC from 2008 to 2010, a Director and Principal Scientist of Consumer Security at PayPal from 2010 to 2013, a Senior Director at Qualcomm from 2013 to 2015, Chief Scientist at Agari from 2016 to 2018, and Chief of Security and Data Analytics at Amber Solutions from 2018 to 2020.

14. Agari is a cybersecurity company that develops and commercializes technology to protect enterprises, their partners and customers from advanced email phishing attacks. At Agari, my research studied and addressed trends in online fraud, especially as related to email, including problems such as Business Email Compromise, Ransomware, and other abuses based on social engineering and identity deception. My work primarily involved identifying trends in fraud and computing before they affected the market, and developing and testing countermeasures, including technological countermeasures, user interaction and education. I made heavy use of virtualization in my work, and published on matters related to virtualization. Example publications include "Detecting computer security risk based on previously observed communications" (U.S. Patent No. 10,715,543 B2, 2016 priority date) and "Using message context to evaluate security of requested data" (U.S. Patent No. 10,805,314 B2, 2017 priority date).

15. Amber Solutions is a cybersecurity company that develops home and office automation technologies. At Amber Solutions, my research addressed privacy, user interfaces and authentication techniques in the context of ubiquitous and wearable computing. My work often related to cloud computing, as exemplified in "Configuration and management of smart nodes with limited user interfaces" (U.S. Patent No. 10,887,447 B2, 2018 priority date).

16. In 2020, following my time at Amber, I joined ByteDance—a multinational

Internet company that specializes in developing a range of creative content platforms and applications, including TikTok and other international products. I was employed as a Principal Scientist at ByteDance until April 2021.

17. I am currently the Chief Technology Officer at ZapFraud, a cybersecurity company that develops techniques to detect deceptive emails, such as Business Email Compromise ("BEC") emails. At ZapFraud, my research studies and addresses computer abuse, including social engineering, malware and privacy intrusions, and has also involved studying security aspects related to cloud computing. My work primarily involves identifying risks, developing protocols and user experiences, and evaluating the security of proposed approaches. Using virtual machines is an important tool for security researchers, as it allows them to study the impact of a given executable, website, or other context on a computational environment. My students, direct reports and I have often relied on this approach in the context of interacting with online criminals, their websites and emails they send, as virtual machines enable us to identify their motives and report them to the authorities without making us easy for them to track or identify. Doing this is important to identify new and trending attacks.

18. I have founded or co-founded several successful computer security companies. In 2005 I co-founded RavenWhite Cybersecurity, a provider of authentication solutions, and I am currently its Chief Technical Officer. RavenWhite owns intellectual property related to cloud computing, such as "Cloud authentication" (U.S. Patent No. 10,348,720, 2006 priority date). In 2007 I co-founded Extricatus, one of the first companies to address consumer security education. In 2009 I founded FatSkunk, a provider of mobile malware detection software; I served as Chief Technical Officer of FatSkunk from 2009 to 2013, when FatSkunk was acquired by Qualcomm and I became a Qualcomm employee. FatSkunk developed anti-virus technology, an area of

research where virtualization is common for purposes of protection of assets. In 2013 I founded ZapFraud—the same company described above, where I currently serve as Chief Technical Officer. In 2014 I co-founded RightQuestion, a security consulting company.

19. I have additionally served as a member of the fraud advisory board at LifeLock (an identity theft protection company); a member of the technical advisory board at CellFony (a mobile security company); a member of the technical advisory board at PopGiro (a user reputation company); a member of the technical advisory board at MobiSocial dba Omlet (a social networking company); and a member of the technical advisory board at Stealth Security (an anti-fraud company). I have provided anti-fraud consulting to KommuneData (a Danish government entity), J.P. Morgan Chase, PayPal, Boku, and Western Union.

20. I have authored seven books and over 100 peer-reviewed publications, and have been a named inventor on over 100 patents and patent applications. A large number of these relate to cloud computing, virtualization, virtual machines, and related topics.

21. Relevant publications of mine related to cloud computing include *Controlling data in the cloud: outsourcing computation without outsourcing control* (Proceedings of the 2009 ACM Workshop on Cloud Computing Security, 85-90), *Authentication in the clouds: a framework and its application to mobile users* (Proceedings of the 2010 ACM Workshop on Cloud Computing Security, 1-6), and *Secure server-aided signature generation* (International Workshop on Public Key Cryptography PKC 2001: Public Key Cryptography 383-401).

22. Relevant patent filings of mine related to cloud computing include "Cloud authentication" (U.S. Patent No. 10,348,720, 2006 priority date), Artifact modification and associated abuse detection" (U.S. Patent Application Publication No. 2020/0053111 A1, 2018 priority date), "Authentication translation" (U.S. Patent No. 10,929,512, 2011 priority date), and

"Secure distributed computation in cryptographic applications" (U.S. Patent No. 6,950,937, 2001 priority date).

23. Publications of mine that relate to virtual machines include *Controlling data in the cloud: outsourcing computation without outsourcing control* (Proceedings of the 2009 ACM Workshop on Cloud Computing Security, 85-90), and *Crimeware: Understanding New Attacks and Defenses* (Wiley, 2008).

24. Patent publications of mine that relied on the use of virtual machines include "Detecting computer security risk based on previously observed communications" (U.S. Patent No. 10,715,543, 2016 priority date), "Using message context to evaluate security of requested data" (U.S. Patent No. 10,805,314, 2017 priority date), "Scam evaluation system" (U.S. Patent Application Publication No. 2020/0067861 A1, 2014 priority date) and "System and method for exacting a system resource access cost" (U.S. Patent No. 7,143,163, 2000 priority date).

25. Much of my work has related to cloud environments and platforms (improving them, securing them against abuse, or using them in general, particularly to protect researchers who interact with criminals), the development of distributed protocols, and the use of virtualization. Above, I have provided some examples of my publications relating to these topics. My work has also included research in the areas of applied security, privacy, cryptographic protocols, authentication, malware, social engineering, usability, and fraud.

26. I have been engaged as a technical expert in over 50 computer-related cases, including numerous cases involving Internet security. I have been admitted and recognized in U.S. District Courts as a technical expert in over 20 separate District Court patent trials, as well as before the Patent Trial and Appeal Board ("PTAB").

III. MATERIALS CONSIDERED

27. In forming the opinions that I express in this declaration, I have reviewed the Asserted Patents, the file histories for the Asserted Patents, the documents to which any of the Asserted Patents claim priority, and all documents I cite in this declaration.

IV. UNDERSTANDING OF GOVERNING LAW

28. I am not an attorney, but have been instructed in and applied the law as described in this section.

A. General Claim Construction Principles

29. I understand that the purpose of claim construction is to give claim terms the meaning understood by a POSITA at the time of the claimed invention, when considered in the context of the patent.

30. I understand that claim terms are interpreted in the context of not only the claim in which the disputed term appears, but in the context of the entire patent, including the specification. I also understand that history of that patent's prosecution before the U.S. Patent & Trademark Office ("PTO") is relevant to the meaning of the claims.

31. To that end, I understand that that claims are not to be interpreted in a vacuum and must be considered within the context of this intrinsic evidence. I understand that this "intrinsic evidence" includes the subject patent itself, the history of that patent's prosecution before the PTO, patents (or other materials) incorporated by reference into the subject patent, and patents (or other materials) cited to or by the PTO during the prosecution of the subject patent. Additionally, I understand that the patentee can act as his or her own "lexicographer" and provide his or her own special definition of a claim term, even if the patentee's definition is different than the ordinary usage of that term. When a patentee clearly acts as his or her own lexicographer, I understand that the definition provided by the patentee should be applied.

32. I further understand that dictionary definitions and other extrinsic evidence, such as expert testimony, may be considered to the extent it aids in the understanding claim terms, but extrinsic evidence may not be used to vary the meaning given to a claim term based on the intrinsic evidence.

B. Construing Claims Under 35 U.S.C. § 112(f)

33. I understand that, according to 35 U.S.C. § 112(f), an element in a claim for a combination may be expressed as a means or step for performing a specified function without the recital of structure, material, or acts in support thereof, and such a claim must be construed to cover only the corresponding structure, material, or acts described in the specification and equivalents thereof.

34. I understand that, with this provision, Congress struck a balance in allowing patentees to express a claim limitation by reciting a function to be performed rather than by reciting structure for performing that function, while placing specific constraints on how such a limitation is to be construed, namely, by restricting the scope of coverage to only the structure, material, or acts described in the specification corresponding to the claimed function, and equivalents thereof.

35. When a claim term lacks the word "means," there is a presumption that that term is not subject to § 112(f). I understand, however, that this is not a strong presumption and can be overcome. The key inquiry in determining whether a claim term is subject to § 112(f) is whether the words of the claim are understood by persons having ordinary skill in the art to have a sufficiently definite meaning as the name for structure. Where the words of the claim fail to recite sufficiently definite structure or else recites function without reciting sufficient structure for performing that function, I understand that the presumption is overcome and the corresponding claim term must be construed under § 112(f) as discussed above.

36. I understand when a limitation subject to § 112(f) claims a functionality that is performed by a computer, the corresponding structure should not be construed as a general purpose computer or a microprocessor, unless the claimed function is one that any processor can perform without special programming. Instead, the corresponding structure must include the algorithm disclosed in the specification for performing the claimed function, and equivalents thereof. I understand that the algorithm can be expressed in a mathematical formula, in prose, in a flow chart, or in any other manner that provides sufficient structure, such that a POSITA would understand the boundaries of the claim. I further understand that if a specification fails to recite, or does not recite a sufficient algorithm where one is required, the claim term at issue is indefinite and therefore invalid.

V. LEVEL OF SKILL IN THE ART

37. I have been informed and understand that the level of ordinary skill in the relevant art at the time of the invention is relevant to inquiries such as the meaning of claim terms.

38. I have been informed and understand that factors that may be considered in determining the level of ordinary skill include: (1) the education of the inventor; (2) the type of problems encountered in the art; (3) prior art solutions to those problems; (4) rapidity with which innovations are made; (5) sophistication of the technology; and (6) education level of active workers in the relevant field. I have been further informed and understand that a person of ordinary skill in the art is also a person of ordinary creativity.

39. A POSITA at the time of the alleged inventions would have had at least a bachelor's degree in computer science, computer engineering, or an equivalent degree and at least three years of work experience in the field of cloud computing, computer networks and systems, and/or a similar field. More direct industry experience can compensate for less formal education, and more formal education can compensate for less industry experience. This

POSITA would have been aware of and generally knowledgeable about the structure and operation of virtual machines and different types of hypervisors.

40. In view of my educational background (*e.g.* a Ph.D. in Computer Science from University of California at San Diego, obtained in 1997) and decades of experience working with computer systems, cloud computing, network security, fraud prevention, and virtual machines, as discussed above, I was a POSITA when the ’730 Patent was filed. Although I was a person of more than the ordinary level of skill in the art as of November 14, 2003, the filing ate of the ’730 Patent, my opinions herein were formed considering, and are expressed from, the perspective of an ordinarily skilled artisan. Further, I was a POSITA when the ’612 Patent was filed. Although I was a person of more than the ordinary level of skill in the art as of April 14, 2010, the filing date of the ’612 Patent, my opinions herein were formed considering, and are expressed from, the perspective of an ordinarily skilled artisan.

41. My opinions in this report are from the perspective of a POSITA. For example, my discussion regarding the ’730 Patent is from the perspective of POSITA in mid-November 2003 or earlier and my discussion regarding the ’612 Patent is from the perspective of POSITA in mid-April 2010 or earlier.

VI. EFFECTIVE FILING DATE OF THE PATENTS

42. I have reviewed the specification of the ’730 Patent. I understand that the application leading to the ’730 Patent was filed on November 14, 2003. For the purposes of this declaration, I have been instructed to use November 14, 2003 as the effective priority date of the ’730 Patent. My opinions in this declaration were formed from the perspective of a person of ordinary skill in the art as of November 14, 2003, including both the knowledge of a person or ordinary skill in the art at that time as well as how a person of ordinary skill in the art would have understood the prior art at that time.

43. I have reviewed the specification of the ’612 Patent. I understand that the application leading to the ’612 Patent was filed on April 14, 2010. For the purposes of this declaration, I have been instructed to use April 14, 2010 as the effective priority date of the ’612 Patent. My opinions in this declaration were formed from the perspective of a person of ordinary skill in the art as of April 14, 2010, including both the knowledge of a person or ordinary skill in the art at that time as well as how a person of ordinary skill in the art would have understood the prior art at that time.

VII. OPINIONS RELATED TO THE ’730 PATENT

A. “resource management logic to distribute server resources to each of the plurality of virtual machines according to current and predicted resource needs of each of the multiple workloads utilizing the server resources”

Microsoft’s Proposed Construction	Daedalus Blue’s Proposed Construction
Under the *Williamson* doctrine, this term is a means-plus-function phrase under 35 U.S.C. §112(f). Function: “to distribute server resources to each of the plurality of virtual machines according to current and predicted resource needs of each of the multiple workloads utilizing the server resources.” Structure: indefinite.	Plain and ordinary meaning, not subject to 35 U.S.C. §112(f).

44. It is my opinion that “resource management logic to distribute server resources to each of the plurality of virtual machines according to current and predicted resource needs of each of the multiple workloads utilizing the server resources” is indefinite because “resource management logic” is a “means-plus-function” term and there is no structure (here, an algorithm) disclosed in the specification that corresponds to the “resource management logic” for performing the claimed functions.

45. Specifically, based on my knowledge and experience, the phrase “resource management logic” would not have been understood by a POSITA as having a definite meaning.

Indeed, this phrase had no recognized meaning to a POSITA at the claimed date of invention of the '730 Patent. "Resource management logic" is not a standard component of a computer system and does not connote any specific structure to a POSITA.

46. The phrase simply recites "logic" to perform the functions of "to distribute server resources to each of the plurality of virtual machines according to current and predicted resource needs of each of the multiple workloads utilizing the server resources" without any structure associated with the function. The terms "resource management" refers simply to what the logic does but has no specific meaning to a POSITA regarding the structure of the "logic."

47. Thus, a POSITA would have understood "resource management logic" to be a functional term, without reciting any definite structure. Therefore, I understand that this term should be governed by §112(f).

48. As noted above, I understand that when a term is governed by § 112(f), a POSITA must look to the specification to determine what structure, if any, is disclosed for performing the claimed function. Based on my review of the specification, a POSITA would have appreciated that the functions would have been performed by "logic" implemented in a computer system. For example, the specification makes it clear that the system of the invention "includes a plurality of servers to be utilized by multiple workloads." '730 Patent, Abstract. The Field of the Invention also states that "[t]he present invention relates to the field of virtual machine (VM) based hosting architectures and, more specifically, to a method of providing improved server utilization in a VM based hosting architecture." *Id.*, 1:6-9. The Summary of the Invention states that "[a]ccording to the present invention … the system includes resource management logic to distribute server resources." *Id.*, 1:37-47. Given that the system itself is made up of servers with VMs, and the specification states that this system of the invention *includes resource management*

logic, a POSITA would have understood that the claimed functions are performed by logic operating on the computing system of the invention.

49. However, it is my opinion that the specification fails to provide an algorithm for perform the claimed function.

50. Upon my review of the specification, I found that the following is the only relevant part of the specification:

> At block 80, the load balancers included within hosting architecture 10 **predict** the resource requirements for each of their respective customers every T seconds (based on customer workload data collected by the load balancers), and send the **predicted** resource requirements to global resource allocator 26. In the exemplary embodiment the resource requirement **predictions** can be **<u>utilizing one of many algorithms</u>**.

Id., 7:4-11 (emphasis added).

51. First, this confirms that the claimed functionality is indeed carried out by an "algorithm" because it states that the "predictions" can "utiliz[e] one of many algorithms." However, a POSITA would have found that while the specification explicitly calls for usage of an "algorithm" at least for the predictions, there is no algorithm disclosed in the specification. Rather, the disclosure merely recites that the load balancers can predict the resource requirements "utilizing one of many algorithms" without identifying any particular algorithms that can be used.

52. After examining the rest of the specification, I have found other discussions of the "prediction," but these sections of the specification fail to shed light on the prediction algorithm. Rather, these sections simply discuss using the prediction information. *See e.g. id.*, 7:4-32, Abstract, 1:37-59, 5:39-56, Claim 1. For example, the Abstract states that the "system includes resource management logic to distribute server resources to each of the plurality of VMs

according to predicted resource needs." *Id*., Abstract. This simply regurgitates the language from Claim 1. The Summary of the Invention contains a similar regurgitation of Claim 1's language. *See id.,* 1:37-59. Thus, it is my opinion that the specification fails to disclose any algorithm that performs the claimed "prediction" function.

53. Regarding the claimed determination of "current … resource needs of each of the multiple workloads utilizing the server resources," my review of the specification shows that the only relevant disclosures are contained at column 5 line 32 to 38 and column 6 line 64 to column 7 line 3. Specifically, the specification states:

> The load balancers (e.g., load balancer 50, load balancer 52 and load balancer 54) for each of the customers will measure the current offered load and send the information to the global resource allocator 26. For example, in hosting architecture 10 load balancer 50 measures the current offered load for customer one 18 and sends that measurement to the global resource allocator.

Id., 5:32-38.

> At block 78, each of the load balancers (e.g., load balancers 50, 52, 54 and 56) in the hosting system measure incoming load for each of their respective customers (e.g., customers 18, 20, 22 and 24) and saves the measured incoming loads in a local database. In an alternative embodiment, a single load balancer may serve multiple customers or all customers in the hosting system.

Id., 6:64-7:3.

54. These disclosures confirm that the claimed function is performed by the computer system, but they reveal no algorithms setting forth how the claimed "logic" actually determines the "current … resource needs of each of the multiple workloads utilizing the server resources." *Id*., 8:9-12. For example, there is no disclosure of how the load balancers measure the current loads, including what algorithms the load balancers may use to determine what the current load is. Furthermore, there is no disclosure of how the resource management logic determines the

current resource needs of any workload, based on the measurements of the incoming load. For example, there is no disclosure of how the logic converts the measured load into any kind of an evaluation or determination of which resources (*e.g.* CPU, memory, network bandwidth, storage space) may be required simply because a "load" has been measured.

55. Accordingly, it is my opinion that the "resource management logic" term is functional, not structural, and that § 112(f) applies. It is also my opinion that the specification fails to disclose an algorithm which corresponds to the "logic" for performing the claimed functions. As a result, it is my opinion that the claim is indefinite.

B. "global resource allocator (GRA) … for receiving said offered workload messages and assigning an optimum matching of combinations of whole integer numbers of workload servers and fractional virtual workload servers that the GRA controls to each of the respective customer workloads according to identified resource requirements"

Microsoft's Proposed Construction	**Daedalus Blue's Proposed Construction**
Under the *Williamson* doctrine, this term is a means-plus-function phrase under 35 U.S.C. §112(f). Function: "receiving said offered workload messages and assigning an optimum matching of combinations of whole integer numbers of workload servers and fractional virtual workload servers that the GRA controls to each of the respective customer workloads according to identified resource requirements" Structure: The structure corresponding to the global resource allocator that performs the claimed function is a computer-implemented algorithm that includes the following steps: 1. Split server resources between VMs evenly to start (*see* '730 Patent, 5:24-56); 2. Receive measurements and/or prediction data from the load balancer(s) (*see id.*, 2:9-19, 2:42-52, 5:24-56, Claim 11, Fig. 3A); 3. Predict what resources are needed by each customer (*see id.*, 5:24-56, Fig. 3A);	Plain and ordinary meaning, not subject to 35 U.S.C. §112(f).

4. Determine if any server capacity would be exhausted based on the predicted resource requirements (*see id.*, 6:58-7:58, Fig. 3A); IF NO GO TO STEP 5; IF YES GO TO STEP 6. 5. (FROM STEP 4: If no), adjust resource allocation for each of the VMs on all servers to conform with the prediction (*see id.*, 5:24-56, 6:58-7:58, Fig. 3A). 6. (FROM STEP 4: If yes), mark the servers as overloaded. Contact the resource control agents at each server with resource assignments for each virtual machine pursuant to a process of moving load from the overloaded servers (*see id.*, 6:58-7:58, Figs. 3A, 3B).	

56. It is my opinion that "global resource allocator (GRA)" is a means-plus-function term under § 112(f) because it is functional language that does not convey structure.

57. Specifically, based on my knowledge and experience, this phrase would not have been understood by a POSITA as having a definite meaning. Indeed, this phrase had no recognized meaning to a POSITA at the claimed date of invention of the '730 Patent. A "global resource allocator" is not a standard component of a computer system, which would connote any specific structure to a POSITA.

58. The phrase simply recites an "allocator" to perform the functions of "receiving said offered workload messages and assigning an optimum matching of combinations of whole integer numbers of workload servers and fractional virtual workload servers that the GRA controls to each of the respective customer workloads according to identified resource requirements" without any structure associated with the function. The term "global resource" refers simply to what the allocator does but has no specific meaning to a POSITA regarding the structure of the "allocator."

59. Thus, a POSITA would have understood "global resource allocator" to be a functional term, without reciting any definite structure. Therefore, I understand that this term should be governed by §112(f).

60. As noted above, I understand that when a term is governed by § 112(f), a POSITA must look to the specification to determine what structure, if any, is disclosed for performing the claimed function. Upon reviewing the specification of the patent, I have determined that a POSITA would have understood that the specification discloses the following six-step algorithm for performing the claimed function, which I summarize below:

61. **First**, the specification states that the "resource allocation for each of the VMs is 1/N of the server to begin with," where N is the number of VMs. '730 Patent, 5:24-31. This shows that the first step is to divide the resources equally based on the number of VMs.

62. **Second**, the specification shows that the global resource allocator receives information and/or prediction data from the load balancer(s). *See id.*, 2:9-19, 2:42-52, 5:32-38, Claim 11, Fig. 3A. Specifically, the specification states that the "system can also include at least one load balancer to measure the current offered load. The global resource allocator determines how to distribute the resources between the plurality of VMs, according to the measurements received from the at least one load balancer." *Id.*, 2:9-19.

63. **Third**, the specification shows that global resource allocator may use the received information to generate a prediction of what resources will be needed by each customer, or it may receive predictions from the load balancers. Specifically, the specification states that "the global resource allocator 26 analyzes the received information and generates a prediction as to what resources will be needed by each of the customers." *See id.*, 5:39-43, Fig. 3A. Further the specification states that "load balancers" may "predict the resource requirements for each of their respective customers every T seconds … and send the predicted resource requirements to global resource allocator." *See id.*, 6:58-7:11.

64. **Fourth**, the global resource allocator determines, based on the information

received, "if any server capacity would be exhausted based on the predicted resource requirements for each of the virtual machines." *Id.*, 6:58-7:18, Fig. 3A. Based on this determination the system has two separate courses of action, depending on if the server resources will not be exhausted (thereupon the algorithm continues to the fifth step) or will be exhausted (thereupon the algorithm continues to the sixth step) as shown in the following paragraphs.

65. **Fifth**, if the determination in step four above is "no," the specification states that the global resource allocator "contacts the resource control agent … on each of the servers … to adjust the resource allocation for each of the virtual machines as per the predictions." *Id.*, 7:18-24; *see also id.*, 5:24-56, Fig. 3A. This allows utilization of more server resources when the servers have yet to make use of all its available resources.

66. **Sixth**, if the determination in step four is "yes," "the server(s) with exhausted capacity will be marked as overloaded" and after a reshuffling of load the "global resource allocator 26 contacts each resource control agent … at each server … and provides them with resource assignments for each of their virtual machine." *Id.*, 7:25-58, Figs. 3A, 3B. Specifically, the specification discloses that these severs with exhausted capacity marked "overloaded" will have load moved from them "to the under loaded servers until [the overloaded servers] are no longer overloaded." *Id.*

67. As can be seen in the algorithm above, step two is related to how the global resource allocator receives information and/or prediction data from the load balancers. Thus, a POSITA would have understood that it is step 2 of the above algorithm that performs the claimed "receiv[es] said offered workload messages."

68. Further, the other steps of the algorithm above are related to how resources are assigned based on the needs of customers for their workloads. Thus, a POSITA would have

understood that it is steps 1 and 3 through 6 of the above algorithm that perform the claimed function of "assigning an optimum matching of combinations of whole integer numbers of workload servers and fractional virtual workload servers that the GRA controls to each of the respective customer workloads according to identified resource requirements."

69. Additional descriptions of the "global resource allocator" found elsewhere in the specifications comport with the above algorithm. Specifically, the Summary of the Invention makes it clear that the "global resource allocator determines how to distribute the resources between the plurality of VMs, according to the measurements received from the at least one load balancer" and sends "instructions" to "local resource control agent[s]." *Id*., 2:9-19; *see also id.*, 2:20-58, 4:51-59. Additionally, other descriptions are similarly high level and similarly consistent with the algorithm provided above. The specification later states that the "global resource allocator [is] for allocating resources among the multiple virtual machines hosted on each of the servers." *Id*., 4:51-54. This is a higher-level description of the behavior encompassed by the disclosed six-step algorithm presented above—it does not conflict with or add to the algorithm I discovered in the specification.

70. Accordingly, it is my opinion that "global resource allocator" is a means-plus-function term and that the corresponding structure in the specification for performing the claimed function is given by the six-step algorithm discussed in the paragraphs above.

VIII. OPINIONS RELATED TO THE '612 PATENT

A. "flagging the instance of a VM" for Autonomic Scaling

Microsoft's Proposed Construction	Daedalus Blue's Proposed Construction
Plain and ordinary meaning, which is "adding a marker to the configuration information for a particular VM occurrence indicating whether that VM occurrence is to be autonomically scaled."	Plain and ordinary meaning.

71. It is my opinion that Defendant's construction of "flagging the instance of a VM" for autonomic scaling, (recited in Claims 1, 3, 6, 8, 11, and 13) as "adding a marker to the configuration information for a particular VM occurrence indicating whether that VM occurrence is to be autonomically scaled" is the correct interpretation of the claim language. I note that Plaintiff proposed a slightly longer term for construction here, namely "flagging the instance of a VM for autonomic scaling." Defendant's proposed construction is the correct interpretation of either of these variants of this term, and is how a POSITA would have understood either of these terms when read in context with the rest of the language of the pertinent claim.

72. A POSITA looking to understand the correct meaning of the term "flagging the instance of a VM" would have been guided first by the language of the claims in which this term appears. I observe that the language of Claim 1 requires "deploying, by the cloud [OS], an instance of a VM, including flagging the instance of a VM for autonomic scaling, including termination." '612 Patent, 15:50-52. A POSITA would have recognized that Claim 1 recites "flagging *the instance* of a VM." *Id.*, 15:51 (emphasis added). This indicates to the POSITA that the information which acts as the flag for a particular VM must be associated with that instance of the VM in some manner. Thus, the information must be specifically for the specific VM, as opposed to being some other type of perhaps more general information about the system as a

whole. Claim 1 further recites that this flagging is performed by the cloud OS. A POSITA would have recognized from the language of Claim 1 that the cloud OS is the same component responsible for assembling and installing the specifications for that VM instance (*i.e.*, deployment). *Id.,* 15:50-51, 16:1-8.

73. A POSITA would have been further guided by the teachings of the specification of the ’612 Patent about how the inventors described the “flagging” process. In particular, I observe that ’612 Patent specification teaches that the “flagging” process is done at the level of individual VM instances. For example, the inventors taught that their idea of “flagging” included “storing” data that was used in the autonomic scaling process. One example the inventors taught was the use of “predetermined threshold values of operating characteristics.” *Id.*, 9:51-52. Another key teaching of the specification, which would have guided a POSITA as to the meaning of the “flagging” limitations, is that the data used in the autonomic scaling process is stored “in the cloud operating system in association with an identifier (204) of the instance of a VM.” *Id.*, 9:46-53. In other words, in deploying a particular VM instance that should be autonomically scaled, the OS stores an indicator in that VM instance’s configuration code to indicate that the VM instance, once deployed, should be subject to scaling. *See e.g.,* Declaration of Markus Jakobsson, Ph.D., Ex. 2 (Rietschote) at 12:1-34; Ex. 3 (Sharma) at 16:38-17:14. The fact that the autonomic scaling data is stored “in association with” an identifier of the instance of the VM teaches a POSITA that this data would need to be “stor[ed]” in the configuration information for a particular VM instance (for example, in the VM specifications), because otherwise, it would not be associated with that particular VM instance.

74. My analysis of the ’612 Patent’s file history confirms that the inventors understood that the “flagging the instance of a VM” term referred to a flag for a particular VM

occurrence. For example, the Applicant, in arguing to overcome the Ferris reference, acknowledged that the portions of the Ferris reference that the Examiner cited for "flagging" did "indicate the resources [ ] available for addition," but argued that they did not indicate their availability for termination. Def's Br., Ex. 7 ('612 Patent File History), 10/10/2012 Applicant Arguments at 12. I observe that Ferris talks about data associated with "each participating server in a set of resource servers 108 indicating a status regarding the provisioning of their respective resources," and was clearly operating on an individual-server basis. Def's Br., Ex. 8, 7/10/2012 Non-Final Rejection at 3. The Applicants agreed that this teaching of Ferris is an indication of resources available for addition, for "each participating server." This indicates to a POSITA that the Applicant agreed that flagging the instance of the VM meant that information (*i.e.*, a marker) indicating that the VM was subject to autonomic scaling was tied to each individual VM, namely by storing that information with the VM configuration data.

75. Similarly, in its Appeal Brief, Applicant argued that the Examiner's reliance on the "state flag" in a reference (Yach) asserted by the Examiner was improper, because that "flag" did not indicate whether the VM is amenable to autonomic scaling including termination. Def's Br., Ex. 9 ('612 Patent File History), 6/4/2013 Appeal Brief at 6-7. Rather, Applicant argued, Yach's flag provided the VM's state—*i.e.*, executing or idle—and was cleared when the VM entered an idle state. *Id*. at 7. A POSITA would have understood that Yach also associated its flags with an individual VM instance. *Id.,* at 6-7. Indeed, the Applicant confirmed in the file history its own understanding of this claim term when it agreed that Yach disclosed "flagging a virtual machine," but argued that Yach's flagging was for a different purpose than taught by the '612 Patent. *Id.*

76. Accordingly, it is my opinion that the "flagging" limitations of Claims 1, 6 and 11

would have been understood to require adding a marker to the configuration information for a particular VM occurrence indicating whether that VM occurrence is to be autonomically scaled.

All statements herein made by me of my own knowledge are true and all statements made on information and belief are believed to be true. I declare under penalty of perjury that the foregoing is true and correct.

Executed this 3rd day of September, 2021 at Portola Valley, California.

Markus Jakobsson, Ph.D.